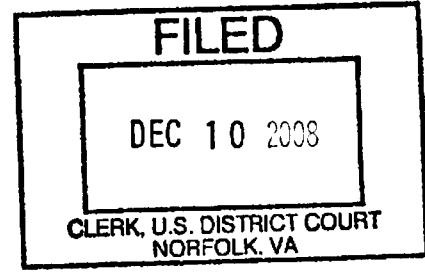


UNITED STATES DISTRICT COURT
EASTERN DISTRICT OF VIRGINIA
Norfolk Division

LEVEL 3 COMMUNICATIONS, LLC,



Plaintiff,

v.

Civil Action No. 2:07cv589

LIMELIGHT NETWORKS, INC.,

Defendant.

OPINION AND ORDER

This matter is before the Court for construction of disputed terms found in claims of three patents-in-suit, as required by *Markman v. Westview Instruments, Inc.*, 517 U.S. 370 (1996). The parties prepared a Joint Statement Regarding Claim Construction (“Joint Statement”), and each side submitted briefs on claim construction issues. The matter is now ripe for decision and the Court considers the merits below.

I. FACTUAL AND PROCEDURAL BACKGROUND

Plaintiff Level 3 Communications, LLC acquired the patents-in-suit, U.S. Patent Numbers 7,054,935 (the ““935 patent”), 6,654,807 (the ““807 patent”), and 6,473,405 (the ““405 patent”), from Savvis Communications Corp. in a deal that was publicly announced in December 2006 and that closed in January 2007. Plaintiff sent Defendant Limelight Networks, Inc. a letter dated February 9, 2007, informing Defendant of Plaintiff’s acquisition of these patents.

On December 17, 2007, Plaintiff filed the instant action against Defendant, alleging infringement of the patents-in-suit. In its Complaint, Plaintiff generally described the technology that is the subject of these patents-in-suit as follows:

A Content Delivery Network (CDN) is a system that supports delivery of information, such as video, music, games, and software, to computer users or computers on behalf of its subscribers (typically content providers). A CDN can have multiple servers distributed at various locations around the U.S. and/or the world.

A content provider such as a website operator can subscribe to a CDN service and then use the CDN for delivery of that content provider's information to computer users or computers. A content provider's use of a CDN is transparent to its end users. For example, when a user requests information/content from a content provider that has subscribed to a CDN service, some or all of that requested information may be delivered to the user from one or more of the CDN's servers, instead of directly from the content provider itself.

A content provider that subscribes to a CDN service may offload substantial overhead and distribution responsibility to the CDN service, and CDN services provide a number of benefits and advantages to both content providers and end users. These include (a) quicker and more efficient delivery of the content providers' information, providing a better experience to their end users; (b) allowing content providers to avoid building and maintaining their own large networks of servers, thereby reducing their costs; (c) reducing load (and potential overload) on the content providers' own servers; and (d) scalability – allowing content providers to support extra capacity when needed.

Complaint ¶¶ 17–19. Plaintiff alleges that there is significant value in using such a shared infrastructure to deliver the content of third parties to end users. Plaintiff alleges that Defendant is directly and/or indirectly infringing one or more claims in each of the aforementioned patents, and seeks to enjoin Defendant's alleged infringement and recover money damages. Defendant generally denies Plaintiff's allegations, and requests award of its costs and attorneys' fees for defending the suit.

On July 14, 2008, a *Markman* hearing was conducted before U.S. District Judge Robert G. Doumar for the purpose of construing the claims.¹ The *Markman* hearing concerned sixteen

¹ Former U.S. District Judge Walter D. Kelley, Jr. presided over this case until February 27, 2008, when the case was reassigned to U.S. District Judge Robert G. Doumar.

terms relating to the '935 and '807 patents and six terms relating to the '405 patent, for a total of twenty-two terms initially in dispute. The parties indicated in their Joint Statement, filed in advance of that hearing, that they agreed on the construction of two claim terms in the '405 patent: "cost" and "existing routing mechanisms." As discussed below in greater detail, in the course of the July 14, 2008 *Markman* hearing, the parties came to agreement on the construction of four additional terms, leaving a total of eighteen terms to be decided by the Court. At the conclusion of the *Markman* hearing, Judge Doumar requested that the parties agree upon an expert in the field with whom the Court could consult. The parties subsequently agreed upon Ellen W. Zegura, D.Sc., who is Professor, Associate Dean, and Chair of the Computing Science and Systems Division of the College of Computing at the Georgia Institute of Technology in Atlanta, Georgia. The Court requested that Dr. Zegura review the disputed claim terms and the Joint Statement, and provide a report with a suggested resolution as to each term. Dr. Zegura did so, and the parties responded with their positions on her suggested definitions as to the disputed claim terms.

On September 17, 2008, before Judge Doumar had made any rulings in connection with the *Markman* hearing, this case was reassigned to U.S. District Judge Jerome B. Friedman. On September 19, 2008, the parties conducted a status conference with Judge Friedman by telephone. For scheduling reasons, the case was reassigned to this Judge on October 6, 2008, and Judge Friedman advised the parties of this by letter that same day. On October 9, 2008, the parties conducted a status conference with the Court, in which the status of outstanding motions and the possibility of the need for a supplemental *Markman* hearing were discussed. With the parties' submissions and the July 14, 2008 hearing transcript before it, the Court advised the

parties on October 16, 2008 that it considered a supplemental *Markman* hearing to be unnecessary for a decision on the claim construction motions. Instead, on October 23, 2008, the Court conducted an on-the-record telephonic status conference, in which the parties answered a handful of technical questions posed by the Court. After careful consideration of the briefs and other materials submitted by the parties, the record before the Court, the argument and discussion of counsel at the July 14, 2008 hearing (as reflected by the transcript of that hearing), and the argument at the October 23, 2008 status conference, the Court issues this Opinion and Order detailing its construction of the disputed claim terms.

II. CLAIM CONSTRUCTION PROCEDURE

In *Markman*, the U.S. Supreme Court succinctly explained the basis for, and importance of, claim construction:

The Constitution empowers Congress “[t]o promote the Progress of Science and useful Arts, by securing for limited Times to Authors and Inventors the exclusive Right to their respective Writings and Discoveries.” Art. I, § 8, cl. 8. Congress first exercised this authority in 1790, when it provided for the issuance of “letters patent,” Act of Apr. 10, 1790, ch. 7, § 1, 1 Stat. 109, which, like their modern counterparts, granted inventors “the right to exclude others from making, using, offering for sale, selling, or importing the patented invention,” in exchange for full disclosure of an invention, H. Schwartz, Patent Law and Practice 1, 33 (2d ed. 1995). It has long been understood that a patent must describe the exact scope of an invention and its manufacture to “secure to [the patentee] all to which he is entitled, [and] to apprise the public of what is still open to them.” *McClain v. Ortmayer*, 141 U.S. 419, 424 (1891). Under the modern American system, these objectives are served by two distinct elements of a patent document. First, it contains a specification describing the invention “in such full, clear, concise, and exact terms as to enable any person skilled in the art . . . to make and use the same.” 35 U.S.C. § 112; see also 3 E. Lipscomb, Walker on Patents §10:1, pp. 183–184 (3d ed. 1985) (Lipscomb) (listing the requirements for a specification). Second, a patent includes one or more “claims,” which “particularly point[] out and distinctly claim[] the subject matter which the applicant regards as his invention.” 35 U.S.C. § 112. “A claim covers and secures a process, a machine, a manufacture, a composition of matter, or a design,

but never the function or result of either, nor the scientific explanation of their operation.” 6 Lipscomb § 21.17, at 315–316. The claim “define[s] the scope of a patent grant,” 3 *id.* § 11:1, at 280, and functions to forbid not only exact copies of an invention, but products that go to “the heart of an invention but avoids the literal language of the claim by making a noncritical change,” Schwartz, *supra*, at 82. . . .

Characteristically, patent lawsuits charge what is known as infringement, Schwartz, *supra*, at 75, and rest on allegations that the defendant “without authority ma[de], use[d] or [sold the] patented invention, within the United States during the term of the patent therefor” 35 U.S.C. § 271(a). Victory in an infringement suit requires a finding that the patent claim “covers the alleged infringer’s product or process,” which in turn necessitates a determination of “what the words in the claim mean.” Schwartz, *supra*, at 80; see also 3 Lipscomb § 11:2, at 288–290.

Markman, 517 U.S. at 373–74 (omitted footnote explaining:

Thus, for example, a claim for a ceiling fan with three blades attached to a solid rod connected to a motor would not only cover fans that take precisely this form, but would also cover a similar fan that includes some additional feature, *e.g.*, such a fan with a cord or switch for turning it on and off, and may cover a product deviating from the core design in some non-critical way, *e.g.*, a three-bladed ceiling fan with blades attached to a hollow rod connected to a motor. H. Schwartz, *Patent Law and Practice* 81–82 (2d ed. 1995)).

It is well-settled that a determination of infringement requires a two-step analysis: “First, the court determines the scope and meaning of the patent claims asserted” and second, “the properly construed claims are compared to the allegedly infringing device.” *Cybor Corp. v. FAS Techs., Inc.*, 138 F.3d 1448, 1454 (Fed. Cir. 1998) (en banc) (citing *Markman v. Westview Instruments, Inc.*, 52 F.3d 967, 976 (Fed. Cir. 1995) (en banc), *aff’d*, 517 U.S. 370 (1996)). “It is a ‘bedrock principle’ of patent law that ‘the claims of a patent define the invention to which the patentee is entitled the right to exclude.’” *Phillips v. AWH Corp.*, 415 F.3d 1303, 1312 (Fed. Cir. 2005) (en banc) (quoting *Innova/Pure Water, Inc. v. Safari Water Filtration Systems, Inc.*, 381 F.3d 1111, 1115 (Fed. Cir. 2004)); *see also Vitronics Corp. v. Conceptronic, Inc.*, 90 F.3d 1576,

1582 (Fed. Cir. 1996) (“[W]e look to the words of the claims themselves . . . to define the scope of the patented invention.”).

A. Claim Construction Principles

The Federal Circuit has repeatedly stated that “the words of a claim ‘are generally given their ordinary and customary meaning,’” and that “the ordinary and customary meaning of a claim term is the meaning that the term would have to a person of ordinary skill in the art in question at the time of the invention.” *Phillips*, 415 F.3d at 1312–13 (quoting *Vitronics*, 90 F.3d at 1582). This provides “an objective baseline from which to begin claim interpretation” and is based upon “the well-settled understanding that inventors are typically persons skilled in the field of the invention and that patents are addressed to and intended to be read by others of skill in the pertinent art.” *Id.* at 1313. As noted by the Federal Circuit:

It is the person of ordinary skill in the field of the invention through whose eyes the claims are construed. Such person is deemed to read the words used in the patent documents with an understanding of their meaning in the field, and to have knowledge of any special meaning and usage in the field. The inventor’s words that are used to describe the invention—the inventor’s lexicography—must be understood and interpreted by the court as they would be understood and interpreted by a person in that field of technology. Thus the court starts the decisionmaking process by reviewing the same resources as would that person, *viz.*, the patent specification and the prosecution history.

Id. (quoting *Multiform Desiccants, Inc. v. Medzam, Ltd.*, 133 F.3d 1473, 1477 (Fed. Cir. 1998)).

However, “[i]n some cases, the ordinary meaning of claim language as understood by a person of skill in the art may be readily apparent even to lay judges, and claim construction in such cases involves little more than the application of the widely accepted meaning of commonly understood words.” *Acumed LLC v. Stryker Corp.*, 483 F.3d 800, 805 (Fed. Cir. 2007) (quoting *Phillips*, 415 F.3d at 1314). Finally, when construing claim terms and phrases, the Court cannot

add or subtract words from the claims or appeal to “abstract policy considerations” to broaden or narrow their scope. *See SmithKline Beecham Corp. v. Apotex Corp.*, 403 F.3d 1331, 1339–40 (Fed. Cir. 2005); *Quantum Corp. v. Rodime, PLC*, 65 F.3d 1577, 1584 (Fed. Cir. 1995) (“[I]t is well settled that no matter how great the temptations of fairness or policy making, courts do not redraft claims.”).

B. Types of Evidence to Be Considered

In determining the meaning of disputed terms or phrases, the Court should first examine the claim and the specification. The Federal Circuit has stated that “the claims themselves provide substantial guidance as to the meaning of particular claim terms,” and “[b]ecause claim terms are normally used consistently throughout the patent, the usage of a term in one claim can often illuminate the meaning of the same term in other claims.” *Phillips*, 415 F.3d at 1314.

The claims, however, “do not stand alone” and “must be read in view of the specification, of which they are a part.” *Id.* at 1315 (quoting *Markman*, 52 F.3d at 979); *see also Vitronics*, 90 F.3d at 1582 (stating that “the specification is always highly relevant to the claim construction analysis. Usually, it is dispositive; it is the single best guide to the meaning of a disputed term.”); *Multiform Desiccants*, 133 F.3d at 1478 (stating that “[t]he best source for understanding a technical term is the specification from which it arose, informed, as needed, by the prosecution history.”). The specification, as required by statute, describes the manner and process of making and using the patented invention, and “[t]hus claims must be construed so as to be consistent with the specification, of which they are a part.” *Merck & Co. v. Teva Pharms. USA, Inc.*, 347 F.3d 1367, 1371 (Fed. Cir. 2003); *see also Markman*, 517 U.S. at 389 (stating that a claim “term can be defined only in a way that comports with the instrument as a whole.”);

Phillips, 415 F.3d at 1316 (stating that “our cases recognize that the specification may reveal a special definition given to a claim term by the patentee that differs from the meaning it would otherwise possess. In such cases, the inventor’s lexicography governs.”).

In addition to the claim and specification, the Court should consider the prosecution history, which consists of the complete record of the proceedings before the Patent and Trademark Office (“PTO”), including the prior art cited during the examination of the patent. *Phillips*, 415 F.3d at 1317 (citing *Autogiro Co. of America v. United States*, 384 F.2d 391, 399 (Ct. Cl. 1967)). The prosecution history “provides evidence of how the PTO and the inventor understood the patent” and “can often inform the meaning of the claim language by demonstrating how the inventor understood the invention and whether the inventor limited the invention in the course of prosecution, making the claim scope narrower than it would otherwise be.” *Phillips*, 415 F.3d at 1317 (citing *Vitronics*, 90 F.3d at 1582–83); *see also Chimie v. PPG Indus., Inc.*, 402 F.3d 1371, 1384 (Fed. Cir. 2005) (stating that the purpose of consulting the prosecution history in claim construction is to exclude any interpretation that was disclaimed during prosecution).

The Court may also examine extrinsic evidence, which includes “all evidence external to the patent and prosecution history, including expert and inventor testimony, dictionaries, and learned treatises.” *Markman*, 52 F.3d at 980. Technical dictionaries may provide the Court with a better understanding of the underlying technology and the way in which one of skill in the art might use the claim terms. *Phillips*, 415 F.3d at 1318; *see also Vitronics*, 90 F.3d at 1584 n.6. Expert testimony can be useful:

to provide background on the technology at issue, to explain how an invention works, to ensure that the court's understanding of the technical aspects of the patent is consistent with that of a person of skill in the art, or to establish that a particular term in the patent or the prior art has a particular meaning in the pertinent field.

Phillips, 415 F.3d at 1318; *see also Pitney Bowes, Inc. v. Hewlett-Packard Co.*, 182 F.3d 1298, 1308–09 (Fed. Cir. 1999). “However, while extrinsic evidence ‘can shed useful light on the relevant art,’ [the Federal Circuit has] explained that it is ‘less significant than the intrinsic record in determining ‘the legally operative meaning of claim language.’’’’ *Phillips*, 415 F.3d at 1317 (citing *C.R. Bard, Inc. v. U.S. Surgical Corp.*, 388 F.3d 858, 862 (Fed. Cir. 2004) (quoting *Vanderlande Indus. Nederland BV v. Int'l Trade Comm'n*, 366 F.3d 1311, 1318 (Fed. Cir. 2004))). Finally, with respect to general usage dictionaries, the Federal Circuit has noted that “[d]ictionaries or comparable sources are often useful to assist in understanding the commonly understood meaning of words and have been used . . . in claim construction,” and that “[a] dictionary definition has the value of being an unbiased source ‘accessible to the public in advance of litigation.’’’ *Phillips*, 415 F.3d at 1322 (quoting *Vitronics*, 90 F.3d at 1585).² However, the Federal Circuit cautions that “‘a general-usage dictionary cannot overcome art-specific evidence of the meaning’ of a claim term,” that “the use of the dictionary may extend

² In *Phillips*, the Federal Circuit expressly discounted the approach taken in *Texas Digital Systems, Inc. v. Telegenix, Inc.*, 308 F.3d 1193 (Fed. Cir. 2002), in which the court placed greater emphasis on dictionary definitions of claim terms. *Phillips*, 415 F.3d at 1319–24 (“Although the concern expressed by the court in *Texas Digital* was valid, the methodology it adopted placed too much reliance on extrinsic sources such as dictionaries, treatises, and encyclopedias and too little on intrinsic sources, in particular the specification and prosecution history.”). The Federal Circuit reaffirmed the approach in *Vitronics*, *Markman*, and *Innova* as the proper approach for district courts to follow in claim construction, but acknowledged that there was “no magic formula” for claim construction, and that a court is not “barred from considering any particular sources . . . as long as those sources are not used to contradict claim meaning that is unambiguous in light of the intrinsic evidence.” *Phillips*, 415 F.3d at 1324.

patent protection beyond what should properly be afforded by the inventor's patent," and "[t]here is no guarantee that a term is used in the same way in a treatise as it would be by the patentee." *Phillips*, 415 F.3d 1322 (quoting *Vanderlande*, 366 F.3d at 1321). Additionally, "different dictionaries may contain somewhat different sets of definitions for the same words. A claim should not rise or fall based upon the preferences of a particular dictionary editor, or the court's independent decision, uninformed by the specification, to rely on one dictionary rather than another." *Id.*

With the foregoing principles in mind, the Court will now examine the patents and the disputed claim terms.

III. ANALYSIS OF THE DISPUTED CLAIM TERMS

At the outset, the Court wishes to note that the briefs and other materials submitted by the parties, the oral arguments at the July 14, 2008 *Markman* hearing, and the technical explanations provided by the parties on the October 23, 2008 status conference call were all extremely helpful in educating the Court about the nature of the technology at issue in this case. Both parties have done a good job, not only of elaborating their respective arguments with respect to specific claim terms, but also of drawing the Court's attention to the broader themes underlying the parties' respective proposed constructions as a whole.

A. Claim Constructions Agreed by the Parties at the July 14, 2008 *Markman* Hearing

As noted above, the parties initially disagreed regarding a total of twenty-two terms, many of which recur throughout one or more of the patents at issue. In the course of the July 14, 2008 *Markman* hearing, the parties agreed on the construction of four of those terms. The parties confirmed this agreement during the October 23, 2008 telephonic status conference hearing, and

indicated that they were unable to come to an agreement on any of the eighteen remaining terms.

The four agreed claim constructions are as follows:

1. *Subscriber*

“An entity that publishes resources via an origin server and is authorized to have requests for such resources served from the repeater server network.”

2. *Resource*

“One or more data files such as text, images, video, audio, and the like.”

3. *URL*

“Uniform Resource Locator, as commonly used in the World Wide Web.”

4. *Default Path*

“The path that a message would travel if there were no overlay network.”

B. Claim Constructions Still in Dispute in the ‘807 and ‘935 Patents

The Court will address each of the terms in the ‘807 and ‘935 patents still in dispute in the order in which those terms are listed in Exhibit A to the parties’ Joint Statement. First, however, the Court will discuss the nature and organization of the patents’ claims themselves.

The ‘807 Patent

The ‘807 patent has a total of 48 claims. Independent Claim 1 describes the basic iteration of the invention. Dependent claims 2–7, which are each dependent on the immediately preceding claim, describe a supplemental feature of the invention that rejects requests for information that is not from subscribers (2); determines if requested information is cached locally at the repeater server and, if not, obtains it from an origin server (3); caches information obtained from origin servers in response to client requests (4); constructs and sends a reply containing the

requested information to the client (5); tracks details of the repeater server’s transactions (6); and enumerates particular details to be tracked (7), respectively. Independent Claim 12 and its dependent claims (13–18) follow this same structure. Independent Claim 12 itself largely tracks the language of independent Claim 1, except that it refers to subscribers publishing resources (instead of information) via origin servers and contemplates the existence of “at least one” origin server (instead of “the origin servers”).

Independent claims 8 and 9 also largely track the language of independent Claim 1, except that they respectively provide that the subscriber verifying mechanism’s determination of whether the information requested by the client is from a subscriber is “based, at least in part, on a name by which the repeater server was addressed” (8; ‘807 patent col. 27:46–47) or “on an origin server name in a Uniform Resource Locator (URL) associated with the client request,” where the document in which the requested information is embedded was served by an origin server (9; *id.* col. 28:12–14). Dependent claim 10, which relates to both independent Claims 8 and 9, contemplates using “at least information in [a] Hypertext Transfer Protocol (HTTP) header of the client request to determine a name by which the repeater server was addressed.” *Id.* col. 28:25–27. Independent claim 11 effectively combines the basic iteration of the invention from independent Claim 1 with the supplemental features described in dependent claims 2–5, and further contemplates the use of the criteria discussed in independent claims 8 and 9.

Independent Claim 19 is a variant of independent Claim 12, but, like independent Claim 8, bases the subscriber verifying mechanism’s determination, “at least in part, on a name by which the repeater server in receipt of the client request was addressed.” *Id.* col. 30:26–28. Its dependent Claim 20, like dependent Claim 10, uses information in an HTTP header to determine

that name. Independent Claim 21 is also a variant of independent Claim 12, but, like independent Claim 9, bases the subscriber verifying mechanism’s determination, “at least in part, on an origin server name in a [URL] used to make the client request.” *Id.* col. 30:60–62.

Independent Claim 22 describes a method “[i]n a computer network” similar to the foregoing independent claims, except that it provides that the repeater selector mechanism bases its selection of a repeater server to handle a client request, “at least in part, on a location of the client in the computer network.” *Id.* col. 31:8–9. Its dependent claims, 23 and 24, respectively base the repeater server’s decision as to which origin server from which to retrieve a requested resource, at least in part, on a name by which the repeater server was addressed (23) or information in an HTTP header (24). Independent Claim 25 is a more detailed variant of independent Claim 22, which also incorporates Claim 23’s basis for selecting the origin server from which to obtain a requested resource that is not locally cached (a name by which the repeater server was addressed).

Independent Claim 26 defines a “server including a processor adapted and programmed” essentially to perform the functions of a repeater server, incorporating independent Claim 23’s origin server selection basis (a name by which the repeater server was addressed). *Id.* col. 32:12–13. Independent Claim 27 is essentially the same as independent Claim 26, except that it instead incorporates Claims 9 and 21’s origin server selection basis (using an origin server name in the URL associated with the client request). Dependent Claims 28–34 essentially enumerate for this the same supplemental features as those described in dependent claims 2–7, 10, 13–18, and 24.

Dependent Claim 35 describes a method for substituting repeater servers in case of repeater server failure. Dependent Claims 36, 37, 40, and 41 variously describe individual criteria used by the Repeater Selector Mechanism to select a repeater server to handle a client request: server load (36), client location (37), both load and client location (40), relative transmission cost (41). Dependent claims 38 and 39 describe possible specific characteristics of the repeater selector mechanism; namely, inclusion of a network map and co-location with an origin server, respectively. Dependent claims 42 and 44 contemplate particular types of documents in which requested resources are embedded; namely, Hypertext Markup Language (“HTML”) or XML documents and video streams, respectively. Dependent claim 43 contemplates requested resources being embedded in a document served by an origin server. Dependent claims 45–48 describe possible specific characteristics of the subscriber verifying mechanism; namely, co-location with a repeater server (45, 47) or taking the form of a table (46, 48).

The ‘935 Patent

The ‘935 patent has a total of 18 claims. Claims 1, 7, 9, 10, and 18 are independent claims. These independent claims all closely track each other’s language, diverging only in a handful of meaningful ways. First, the independent claims diverge on the issue of whether at least “one” or “some” repeater servers replicate information from at least “one” or “some” of the origin servers; independent Claims 1, 7, 9, and 10 appear to cover all four possible combinations of “one” and “some” in this context. Second, the independent claims variously—and, in a certain sense, interchangeably—characterize the “stuff” requested by the client as a “resource” (Claims 1 and 18), “information” (Claim 7), or “content” (Claims 9 and 10). “Information” is also used in

all of the '935 patent's independent claims to characterize the "stuff" that one or more repeater servers replicate from one or more origin servers (which, in turn, is presumably the same "stuff" requested by clients). Claim 10 discusses substantially the same process outlined in the other independent claims, but does so by describing "at least one repeater server [being] adapted and programmed to" perform that process. Claim 18 contemplates "assigning a repeater server more than one name or address" and then "using at least the name or address" by which the client request addressed the repeater server "in order to ascertain whether the request is for a resource from a subscriber to the repeater server network." Independent claims 7 and 9 only contemplate using a name, not an address, in this connection, and independent claims 1 and 10 do not explicitly mention the use of either in connection with this subscriber verification step. Instead, dependent Claims 3 and 14 detail the use of name and address, respectively, in connection with this aspect of independent Claim 1, and dependent Claim 4 details the use of an origin server name within the URL in this connection. Claims 5 and 8, respectively dependent on Claims 3 and 7, detail using information in an HTTP header to determine the repeater server's name in this same connection. The remaining dependent claims are essentially variations on these themes.

The Disputed Claim Terms in the '807 and '935 Patents

1. *Origin Server*

a. Proposed Definitions

Plaintiff: "A server containing subscriber resources, which can be distributed to one or more repeater servers."

Defendant: "Server from which subscriber resources originate, and to which a client first makes a request for a particular resource."

b. Discussion

During the July 14, 2008 *Markman* hearing, the parties indicated that they agreed on the first half of Defendant's proposed definition: a "server from which subscriber resources originate." The parties confirmed this agreement on the October 23, 2008 conference call. Accordingly, the Court will adopt that portion of Defendant's proposed definition and confine its analysis of this term to the latter portions of the parties' proposed definitions. The parties' remaining dispute in connection with this term relates to the phrase "and to which a client first makes a request for a particular resource" in Defendant's proposed definition.

This dispute has multiple aspects, but seems to focus in large part on a semantic point regarding use of the word "first." Specifically, the parties seem to dispute whether "first" is intended to refer to a *particular* client's first request for a particular resource or, instead, to the first time that *any* client requests a particular resource. It seems clear to the Court that the first time that a resource is requested by *any* client, that resource will ultimately have to come from the origin server, whether the client request is being serviced directly by the origin server or by a repeater server, which itself obtains the resource from the origin server.

In this invention, subscribers, by definition, only publish resources to origin servers, not to repeater servers. Repeater servers do not regularly, systematically replicate resources from origin servers; instead, they obtain resources from origin servers in response to client requests, on a request-by-request basis (and, in certain configurations, cache the resource to service future requests for that same resource). Consequently, the only way that a repeater server can service a client request without obtaining the requested resource from an origin server is if it or another repeater server in its network has already cached that resource in connection with servicing a

prior request for it. Simply as a matter of logic, there cannot have been any such prior requests the *very first time* that *anyone* requests a particular resource, and therefore that request will have to be serviced, directly or indirectly, by the origin server. This sense of “first” is a meaningful and correct one, and the Court believes that it should be reflected in the term’s definition.

There is also a second, distinct sense of “first” involved here: the first time that a *particular* client requests a resource. In this sense, that client might be the first client ever to request a particular resource or the millionth client to request it. In the former situation, the above analysis would apply, and the resource would ultimately have to be supplied by the origin server. In the latter situation, however, one or more repeater servers might already have cached the requested resource in connection with servicing prior requests for the resource by *other* clients. In that case, that particular client’s request would not necessarily have to involve the origin server in any way: it could be directed to the repeater selector mechanism, which reflects it to a repeater server, which provides the requested resource from its own cache or obtains it from another repeater server’s cache, without having to request or obtain the resource from the origin server.³

There is yet another sense of “first” in Defendant’s proposed definition: the concept that, in this invention, the *initial* destination of the client request for a resource is always an origin server, prior to interception and possible reflection by a repeater selector mechanism. However, this limitation is expressly contradicted by the language of both patents’ specifications. Although

³ The Court also notes, by way of extending the analogy, the possibility that, if it is not the first time that the particular client has requested a resource, then the resource might even be cached locally on the client’s own computer, in which case that client’s subsequent request(s) for the resource could be serviced locally by the client’s own computer, without even connecting to any network.

the '807 patent's specification, while discussing "one aspect" of the invention, admittedly states, "First a client makes a request for a particular resource from an origin server" ('807 patent col. 3:1–4, emphasis added), in discussing the operational flow charts of the inventions, it also states "Note that the bottom row of FIG. 2 refers to an origin server, or a reflector, or a repeater, depending on what the URL" input by the client into the client's browser identifies. '807 patent col. 11:9–11. The '935 patent's specification contains exactly the same language. *See* '935 patent col. 11:11–13. This language clarifies that client requests need not in every case be directed to an origin server even as an initial matter, and that other configurations are possible. Accordingly, the Court rejects this aspect of Defendant's argument and will not construe the definition of origin server as supporting it.

On the basis of the foregoing, the Court's definition combines the portion of Defendant's proposed definition, upon which the parties agreed at the July 14, 2008 *Markman* hearing, with a variation on the remainder of Defendant's proposed definition, which makes clear that only the very first request by *any* client for a given resource will *always* ultimately involve obtaining the requested resource from an origin server.

c. Definition

"Server from which subscriber resources originate and to which the first request for a particular resource is ultimately made."

2. Repeater Server

a. Proposed Definitions

Plaintiff: "A server which may receive resources from one or more origin servers and which is used to service resource requests."

Defendant: “Server that replicates some or all of the subscriber resources from one or more origin servers, and that receives client requests directed to the server by a repeater selector mechanism.”

b. Discussion

The parties’ dispute over this term is twofold. First, the parties differ on whether the definition must contain explicit reference to the repeater server’s replication of resources from one or more origin servers. Second, Defendant’s definition seeks to make explicit the fact that repeater servers receive client requests via the repeater selector mechanism.

With respect to the issue of replication, the Court notes that this concept—at least some repeater servers replicating at least some of the information from at least some of the origin servers—is already separately contemplated by the text of every independent claim in both the ‘807 and ‘935 patents. Accordingly, it would be superfluous to include language relating to replication in the definition of this particular term. Indeed, it would actually be incorrect, to the extent that the claim language in the ‘807 and ‘935 patents only requires *one or some* of the repeater servers to replicate content from origin servers. Adopting Defendant’s definition would effectively impose that requirement on *all* repeater servers, contradicting the explicit intent of the claims in a fashion unsupported by the specification text or any other relevant source.

With respect to the issue of repeater servers receiving client requests exclusively via the repeater selector mechanism, the analysis is much the same. Even if the Court were to agree that this is the case, there is simply no need to include such details in the definition of this term, because every independent claim in both the ‘807 and ‘935 patents separately discusses the presence and role of the repeater selector mechanism. Accordingly, even though Defendant’s

contention in this connection might well be correct, the Court need not adopt Defendant's proposed definition. The Court instead finds Plaintiff's definition to be adequate and accurate, and adopts it.

c. Definition

"A server which may receive resources from one or more origin servers and which is used to service resource requests."

3. *Repeater Server Network*

a. Proposed Definitions

Plaintiff: "A network including repeater servers."

Defendant: "One or more repeater servers connected to a master repeater."

b. Discussion

The parties' dispute over this term is relatively straightforward, focusing on whether or not a repeater server network, by definition, must include a master repeater. The concept of the "master repeater" finds no mention in the claims of either patent, but is instead discussed as a feature of a preferred embodiment in the specification text of both patents. Although the specification admittedly contains considerable discussion of the preferred role played by master repeaters in the invention, Plaintiff correctly points out that "neither the claims, the specification, nor the file history require (through disavowal or otherwise) master repeaters to be connected to non-master repeaters." Level 3's Memorandum of Law in Reply to Limelight's Claim Construction Brief ("Reply Mem.") at 8–9. Accordingly, the Court finds no reason to import this particular feature of a preferred embodiment into the definition of what is otherwise a straightforward, relatively non-technical term.

Although Defendant is correct to point out that each term in a claim must be given meaning and to highlight the fact that Plaintiff's proposed definition does little more than rearrange the constituent words in the claim term, the Court believes that nothing more is needed with respect to this term. Although Defendant further argues that the simplicity of Plaintiff's definition would render redundant the separate "plurality of repeater servers" language found throughout the '807 and '935 patents, that argument ignores other formulations in the patents' claims that clearly equate the two. For example, independent Claims 8 and 11 of the '807 patent discuss "the plurality of repeater servers *forming* the repeater server network." '807 patent cols. 27:40–41, 28:46–47 (emphasis added). Independent Claim 22 of the '807 patent likewise discusses "a plurality of repeater servers distinct from the plurality of origin servers and *forming* at least one repeater server network." *Id.* col. 31:5–7 (emphasis added).

The language of the '935 patent's claims is even more explicit on this point. Independent Claim 1 of the '935 patent discusses "a repeater server network *comprising* a plurality of repeater servers." '935 patent col. 25:57–58 (emphasis added). Independent claim 7 of the '935 patent discusses "a plurality of repeater servers *forming* at least one repeater server network;" independent Claim 10 has language to the same effect. *Id.* col. 26:26–27, 63–64 (emphasis added). Independent claim 9 of the '935 patent discusses "a plurality of repeater servers *making up* at least one repeater server network." *Id.* col. 26:45–46 (emphasis added).

Although it seems clear from the claim language of these patents that repeater server networks also share other characteristics—e.g., having subscribers who publish resources to origin servers that are connected to the repeater server network—Defendant has not argued for inclusion of those characteristics in its proposed definition. Moreover, the foregoing citations

demonstrate that a repeater server network, at bottom, is nothing more than the sum of its parts.

Accordingly, Plaintiff's proposed definition is accurate and will be adopted by the Court.

c. Definition

“A network including repeater servers.”

4. *Name (By Which (a) Repeater Server (Is) Addressed)*

a. Proposed Definitions

Plaintiff: “An identifier.”

Defendant: “A word or term that is uniquely assigned to the selected repeater server.”

b. Discussion

The dispute over this term relates to whether or not the name by which a repeater server is addressed must be unique. The parties' arguments in this connection, however, appear to talk past each other, so to speak. For example, Plaintiff correctly asserts that such uniqueness is not explicitly discussed in the claims or specifications of the '807 and '935 patents. Defendant rightly points out that the contested term is not merely the word “name,” but instead “name” as used in the context of the phrase *name by which a repeater server is addressed* or variants thereof and, accordingly, Plaintiff's argument regarding repeater servers sharing domain names, for example, is inapposite.

It is clear from the '935 patent's claims and specification that a particular repeater server can be assigned more than one name or alias. *See, e.g.*, '935 patent cols. 9:65–67, 26:6–8, 32, 52, 27:22–23. However, it is not clear whether the opposite is true—namely, whether a particular name or alias can be assigned to more than one repeater server. The most relevant

discussion of this seems to be in the ‘935 patent’s specification, appearing in the course of discussing an “alternate approach” to implementing the invention.

For instance, if www.example.com is the origin server, names for three repeaters might be created:

wr1.example.com
wr2.example.com
wr3.example.com

The name “wr1.example.com” would be an alias for repeater 1, which might also be known by other names such as “wr1.anotherExample.com” and “wr1.example.edu”.

... For instance, if repeater 1 is addressed as wr1.example.com, then the origin server is “www.example.com”; if it is addressed as “wr1.anotherExample.com”, then the origin server is “www.anotherExample.com”.

‘935 patent cols. 9:60–10:8. This discussion seems to *assume* Defendant’s contention that each alias is uniquely related to a particular repeater server, even though that repeater server can have multiple aliases associated with different origin servers and subscribers. However, it does not address the question head on.

Fortunately, this apparent gap in the patent text does not create a problem here, because the Court does not believe this particular aspect of the dispute to be one appropriately fought or resolved as a matter of claim construction. Even if Defendant’s contention in this connection is correct, it would not be accurate or appropriate to read that limitation into the definition of the term “name,” which, as used in the patent text, by itself means nothing more than “an identifier.”

Thus, although the Court has accordingly decided to adopt Plaintiff’s proposed construction, this decision does not preclude Defendant from arguing its point on the basis of the context in which the term “name” is used in the claims and the specifications of the ‘807 and ‘935 patents.

c. Definition

“An identifier.”

5. Rejecting the Client Request

a. Proposed Definitions

Plaintiff: “Refusing to return a requested resource to a client.”

Defendant: “Returning a reply that rejects a client request when it is determined that the client request is for a resource that is not from a subscriber, based on the origin server identified in the client-request URL.”

b. Discussion

Defendant’s definition is overbroad, because it defines more than the term itself, and largely superfluous, because the term always appears in the specification text in tandem with additional language to the effect that the reason for rejection is that the requested resource is not from a subscriber. Accordingly, it is unnecessary for the definition of this term itself to contain that reason. The fact that the patent examiner noted this reason for rejection as a particular novelty of the invention is not to the contrary for the same reason: the basis for rejection is already separately articulated in tandem with this term. The Court’s decision not to adopt Defendant’s proposed definition will therefore not preclude Defendant from further arguing its point in this connection.

The Court now turns to the definitions proposed by Plaintiff and the Court’s expert, Dr. Zegura. The fundamental difference between these definitions is whether or not “rejecting,” as used in the ‘807 and ‘935 patents, should be construed to mean only the active voice or instead

should be construed to encompass both active and passive voices. As Dr. Zegura noted, “‘Reject’ is generally used in computer networking to mean [an] active” response.

Defendant is correct that Plaintiff’s definition is arguably overbroad, and Dr. Zegura is correct to identify that overbreadth as encompassing both active and passive senses of “rejecting.” Under Plaintiff’s definition, both an actual rejection message sent to the client, and a failure to respond to the client’s request at all, would be covered by the patent. The Court agrees with Dr. Zegura that only the active sense is appropriate in this definition. The phrase, in context, implies a rejection “event,” not a non-event, such as a mere failure to respond to the client’s request. Accordingly, the Court will adopt Dr. Zegura’s proposed definition.

c. **Definition**

“Returning a reply to the client that declines to provide the requested resource.”

6. *Client Request for a Resource*

a. **Proposed Definitions**

Plaintiff: “Message from a client requesting a resource, which is located in a computer network such as the Internet.”

Defendant: “URL that includes the name of the server from which the resource is requested.”

b. **Discussion**

Plaintiff argues that its proposed definition is simply the plain and ordinary meaning of the terms. Defendant’s proposed definition, on the other hand, seeks to incorporate an alleged limitation from the specifications of the ‘807 and ‘935 patents.

Defendant's argument in this connection is simply not compelling. In fact, it is frankly difficult to see how the portion of the specification cited by Defendant in support of its proposed definition ('935 patent col. 10:23–27, 34–35) could possibly be construed to relate to many of the uses of the term at issue here. Defendant seems to be taking a detail specific to a particular step in the invention's process (namely, a repeater server's initial analysis of a client request) and seeking to apply that detail as a limitation to a term that is used throughout the '807 and '935 patents in a variety of contexts.

A client request, as used throughout the patent, is nothing more than the sum of its constituent words. It does not first become a client request when it is intercepted and redirected by a reflector; it is a client request when it is generated by the client. Consequently, Defendant's proposed definition is inaccurate with respect to many uses of the term, and cannot be adopted. Instead, the Court agrees with Plaintiff that this phrase has a plain and ordinary meaning, and adopts Plaintiff's proposed definition.

c. **Definition**

"Message from a client requesting a resource, which is located in a computer network such as the Internet."

7. *Appropriate Repeater Server*

a. **Proposed Definitions**

Plaintiff: "One or more repeater servers which are not too heavily loaded or which are selected by some measure of network cost."

Defendant: “A repeater server determined via the Best Repeater Selector algorithm, based at least upon repeater load in a Load Table aggregated by the master repeater and based upon at least the network distance from the client.”

b. Discussion

Plaintiff’s proposed definition is taken almost verbatim from the specification. *See* ‘807 patent col. 11:20–22. Defendant argues that the criteria enumerated in Plaintiff’s proposed definition are insufficiently clear to guide a jury, and that Defendant’s proposed definition, which incorporates what Defendant claims is “the only disclosed structure by which the appropriate repeater server is selected,” is clearer and more accurate. Defendant Limelight Networks, Inc.’s Claim Construction Brief for U.S. Patent Nos. 6,654,807; 7,054,935; and 6,473,405 (“Def. Mem.”) at 10–11. Moreover, the difference between Plaintiff’s proposed definition (replacing “and” with “or”) and the text of the specification, however slight, is nevertheless meaningful.

The Court’s definition combines the structure of the definition proposed by the Court’s expert, Dr. Zegura, with both the terminology of Plaintiff’s proposed definition and the additional concept embodied in Defendant’s proposed definition. The Court’s definition addresses Plaintiff’s legitimate concern that Dr. Zegura’s definition introduced new terms (e.g., “network performance” and “network topology”) that do not appear in the specification of the patent and would likely confuse a jury. Although Defendant is correct to point out that Plaintiff’s proposed definition is an expansion on the language in the specification (essentially replacing an “and” with an “or”), Plaintiff is also correct to note that, under the principle of claim differentiation, claims 36, 37, 40, and 41 of the ‘807 patent would be superfluous if the measurement criteria described separately in each of those claims were all already encompassed

together in the patent's primary claim. As the Federal Circuit has explained, the principle of claim differentiation counsels that "the presence of a dependent claim that adds a particular limitation gives rise to a presumption that the limitation in question is not present in the independent claim." *Phillips*, 415 F.3d at 1315. Accordingly, in this respect, at least, Plaintiff's more expansive definition is appropriate here.

c. **Definition**

"A repeater server from the plurality of repeater servers selected using some measure of network cost, the load on at least some of the repeater servers, or the location of the client sending the client request."

8. *Repeater Selector Mechanism*

a. **Proposed Definitions**

Plaintiff: "A mechanism which is constructed and adapted to identify an appropriate repeater server from a network of one or more repeater servers for a particular client request."

Defendant: (Applying Means-Plus-Function under 35 U.S.C. § 112, ¶ 6)

Function: "To identify, for a particular request, an appropriate repeater server to handle the request."

Structure: "A software program, co-located with an origin server, that intercepts a request to such origin server, determines via the Best Repeater Selector algorithm which repeater to direct the client request to (based at least upon repeater load in a Load Table aggregated by the master repeater from information sent by repeater servers, and distributed by the master repeater to the repeater selector mechanism, and based upon at least the network distance from the client),

and then provides to the client a modified resource identifier designating the identified repeater server.”

b. Discussion

Plaintiff’s proposed definition of this term is taken directly from the language of the ‘807 patent’s claims. Defendant, however, argues that the patent does not adequately recite the structure underlying this term, and that the Court should therefore construe this term to be “means-plus-function” language within the meaning of 35 U.S.C. § 112, ¶ 6 (“§ 112, ¶ 6”), which provides that:

[a]n element in a claim for a combination may be expressed as a means or step for performing a specified function without the recital of structure, material, or acts in support thereof, and such claim shall be construed to cover the corresponding structure, material, or acts described in the specification and equivalents thereof.

35 U.S.C. § 112, ¶ 6.

The Court starts with the basic presumption that the absence of the term “means” creates a rebuttable presumption that § 112, ¶ 6 does *not* apply. *See Phillips*, 415 F.3d at 1311 (citing *Personalized Media Commc’ns, LLC v. Int’l Trade Comm’n*, 161 F.3d 696, 703–04 (Fed. Cir. 1998)). However, the Federal Circuit has in some circumstances equated certain other words with “means” for purposes of determining the applicability of § 112, ¶ 6. *See MIT v. Abacus Software*, 462 F.3d 1344, 1354 (Fed. Cir. 2006) (“The generic terms ‘mechanism,’ ‘means,’ ‘element,’ and ‘device,’ typically do not connote sufficiently definite structure” and “the term ‘mechanism’ standing alone connotes no more structure than the term ‘means.’”).

Use of these terms, of course, does not in every case inevitably lead to the application of § 112, ¶ 6. The Federal Circuit has noted that “[c]laim language that further defines a generic

term like ‘mechanism’ can sometimes add sufficient structure to avoid 112 ¶ 6.” *Id.* (citing *Greenberg v. Ethicon Endo-Surgery, Inc.*, 91 F.3d 1580 (Fed. Cir. 1996)). However, if such modifying language “is not defined in the specification and has no dictionary definition, and there is no suggestion that it has a generally understood meaning in the art,” then the term will not be held to “connote sufficient structure to a person of ordinary skill in the art to avoid 112 ¶ 6 treatment.” *Id.*

The ‘807 and ‘935 patents studiously avoid using the term “means” in this connection. Instead, this and other claim terms substitute the word “mechanism.” Obviously patentees cannot in all cases avoid application of § 112, ¶ 6 merely by replacing the word “means” with the word “mechanism.” Furthermore, since the concept of a “repeater server” was among the innovations of this invention, it would be difficult to suggest that this term had a generally understood meaning in the art at the time of the invention. Indeed, if there were no additional details about the repeater selector mechanism provided in the claims and the specifications of these patents, the Court might well be inclined to agree with Defendant that § 112, ¶ 6 treatment would be appropriate. However, that is simply not the case here.

The ‘807 patent’s specification and claims discuss at great length the manner in which the repeater selector mechanism software functions. *See* ‘807 patent cols. 11:12–16:49, 33:19–34:9; *see also* ‘935 patent cols. 11:14–16:39. Moreover, the context in which this term appears in the claims includes other terms (such as “appropriate repeater server” and “handling” a client request) that, when combined, provide sufficient structure to avoid § 112, ¶ 6 treatment of this term. In light of this, Plaintiff’s proposed definition is both accurate and sufficient, and the Court will adopt it.

c. **Definition**

“A mechanism which is constructed and adapted to identify an appropriate repeater server from a network of one or more repeater servers for a particular client request.”

9. *Subscriber Verifying Mechanism*

a. **Proposed Definitions**

Plaintiff: “A mechanism, such as a table, used to associate a subscriber with a resource requested by a client.”

Defendant: (Applying Means-Plus-Function under 35 U.S.C. § 112, ¶ 6)

Function: “To verify whether an entity is a subscriber to the repeater server network.”

Structure: “A Subscriber Table, located at the master repeater and propagated to all repeater servers in the network, having information necessary to determine whether the origin server identified in a client-request URL belongs to a known subscriber.”

b. **Discussion**

The parties’ arguments with respect to this term largely track the arguments regarding the claim term “repeater selector mechanism.” Correspondingly, the Court’s analysis is essentially the same, as well. Here, as with “repeater selector mechanism,” the patentee has substituted the word “mechanism” in lieu of the word “means.” However, the ‘807 patent’s claims and specification likewise contain detailed discussions of the manner in which the subscriber verifying mechanism functions. *See, e.g.*, ‘807 patent cols. 10:30–42, 34:20–31. Accordingly, the Court does not consider the term to fall within the scope of § 112, ¶ 6.

Plaintiff correctly points out that Defendant’s proposed definition, in any case, would be invalid under the principle of claim differentiation. Here, dependent Claims 45 and 47 explicitly

contemplate the subscriber verifying mechanism being located at a repeater server, not only at the master repeater. *See* '807 patent col. 34:20–22, 26–28. Furthermore, dependent Claims 46 and 48 specifically describe the subscriber verifying mechanism as a table. As previously noted, “the presence of a dependent claim that adds a particular limitation gives rise to a presumption that the limitation in question is not present in the independent claim.” *Phillips*, 415 F.3d at 1315. Limiting the definition of the subscriber verifying mechanism to the form of a table would render these dependent claims superfluous. Accordingly, the Court will not adopt these aspects of Defendant’s proposed definition.

With respect to Plaintiff’s proposed definition, the Court notes that it diverges textually from the proposed definition of the other “mechanism” term at issue here. The Court sees no reason to omit the “constructed and adapted” clause from this definition. Accordingly, the Court will adopt a definition combining this claim language with Plaintiff’s proposed definition.

c. Definition

“A mechanism, such as a table that associates a resource requested by a client with a subscriber, which is constructed and adapted to verify whether an entity is any one of the plurality of subscribers to the repeater server network.”

10. Embedded

a. Proposed Definitions

Plaintiff: No proposed construction; plain and ordinary meaning.

Defendant: “Resource automatically downloaded to a user’s computer when a browser loads the page.”

b. Discussion

The parties' dispute over the definition of this term is interlinked with their former dispute over the term "resource," which was resolved by the parties in favor of Plaintiff's proposed definition. Although that fact is not entirely dispositive of the ongoing dispute over this term, to the extent the arguments relating to each are interdependent, it does color the analysis. Defendant's proposed definition seeks to distinguish *embedded* resources, which Defendant argues are downloaded automatically at the time the client's browser loads the web page requested by the client, from *links* in a web page, which do not actually contain the resource with which they are associated, but instead merely prompt the download of the associated resource *when a client separately and explicitly selects* (i.e., "clicks on") *them*. Defendant cites the prosecution history in support of this argument. Def. Mem. Exh. 17 (Application No. 09/612,598 File History Amendment dated October 11, 2001) at 16. However, Plaintiff does not argue this point; in fact, Plaintiff acknowledged in its opening claim construction brief that a "link is not a resource itself but, rather, is associated with a resource." Plaintiff's Memorandum of Law in Support of Level 3's Proposed Claim Construction ("Opening Mem.") at 15. The Court agrees with Plaintiff that the plain and ordinary meaning of the word "embedded" is appropriate here, and therefore no construction is needed. This determination, of course, in no way precludes Defendant from drawing the distinction between embedded resources and links; indeed, as noted above, Plaintiff does not appear to contest this point.

c. Definition

No construction is needed; "embedded" has a plain and ordinary meaning.

11. *(Request for a Resource May Be) Handled (by the Repeater Server Network)*

a. Proposed Definitions

Plaintiff: No proposed construction; plain and ordinary meaning.

Defendant: “A selected repeater server handles a request for a resource to an origin server when that request is intercepted and directed to the repeater server by a Best Repeater Selector algorithm co-located with the origin server.”

b. Discussion

As “client request for a resource” and “repeater server network” are separate terms in dispute in this litigation, and have already been discussed and construed by the Court above, the actual disputed term at issue here is the word “handled” as used in conjunction with those terms.

Defendant’s proposed construction is derived directly from the ‘935 patent’s specification, which describes in detail the process by which client requests are “handled” by the repeater server network. Plaintiff argues without elaboration that no construction of “handled” is necessary.

Although Plaintiff’s suggestion that no construction is necessary has a certain immediate, common-sense appeal, further consideration of the ‘935 patent’s claims and specification, taken as a whole, have convinced this Court that more is needed. However, although the thrust of Defendant’s proposed construction is correct, the Court is not inclined to adopt that construction wholesale, either. For one thing, Defendant’s proposed definition is arguably recursive: it contains a slightly reorganized version of the very phrase it purports to define. Moreover, certain aspects of Defendant’s definition are belied by language in the specification. For example, Defendant’s definition requires the Best Repeater Selector algorithm (a term from the patent that itself would likely require further elaboration for a jury to understand it) to be co-located with the

origin server. However, although the specification admittedly states that the reflector is “typically co-located” with the origin server, it later indicates that “it is possible to leave the origin server’s network address as it is and to let the reflector run at a different address or on a different port. In this way the reflector does not intercept requests sent to the origin server, but can still be sent requests addressed specifically to the reflector.” ‘935 patent col. 7:54–59. This language, combined with the fact that the claims make no mention of “co-location,” suggests that this aspect of Defendant’s proposed definition is unnecessary.

The Court has, instead, decided to adopt its own definition. The Court recognizes the rather bold character of its chosen definition, and Plaintiff might well claim that, in adopting it, the Court is “reading too much into” a single word. However, it is the Court’s belief that, absent such a definition, the claims of the ‘935 patent would be impermissibly overbroad. At the heart of this invention lies the reflector mechanism, which intercepts and selectively re-routes client requests for resources among origin and repeater servers. This is no mere preferred embodiment; it is the essence of the invention, or at least an important constituent element of it. This mechanism, however, finds *no explicit mention in the ‘935 patent’s claims, at all*. Thus, if the Court’s definitions of the claim terms in this patent, read in conjunction with one another, did not somewhere incorporate this constituent element of the invention, then the technology protected by the claims of the ‘935 patent would encompass far more than the actual invention described (in great detail and specificity) by the specification. The Court believes that “handled,” being the verb used in both the claims and the specification to denote the processing of a client request by the invention, is the most appropriate term to carry this descriptive burden, so to speak.

c. **Definition (“Handled”)**

“Intercepted or otherwise sent to the reflector, which analyzes the client request, determines whether to reflect the request to a repeater server or to send the request to an origin server, and, if it determines that it will reflect the request to a repeater server, selects the best repeater server for the request using at least some measure of the load on at least some of the repeater servers and some measure of network distance between the location(s) of at least some of the repeater servers and the location of the client sending the request.”

12. *Obtaining a Client Request*

a. **Proposed Definitions**

Plaintiff: No proposed construction; plain and ordinary meaning.

Defendant: “Receiving a client request directed by the repeater selector mechanism to a repeater server.”

b. **Discussion**

As “client request” is a separate term in dispute in this litigation, and has already been discussed and construed by the Court above, the only remaining dispute over this term obviously relates to the word “obtain[ing].” Defendant’s definition seeks to clarify that the *only* way a repeater server obtains a client request is when that request is reflected to the repeater server by a repeater selector mechanism. It is noteworthy here that the only source Defendant cites in support of its proposed definition is the specification of *another* patent: the ‘935 patent. Defendant does not cite to the specification of the ‘807 patent, its prosecution, history or any

other related source. Plaintiff argues without elaboration that no construction of “obtain[ing]” is necessary.

Adoption of Defendant’s proposed definition is unnecessary for at least three reasons. First, the Court agrees with Plaintiff that “obtain[ing],” as used in the ‘807 patent at issue, does not have any technical or otherwise special meaning, and, as such, needs no definition. Second, Defendant’s proposed definition effectively seeks to limit “obtaining” to its use in a more specific description in independent Claims 8 and 11, both of which specifically contemplate a repeater server obtaining a client request via the repeater selector mechanism. Independent Claim 1, however, contains no such specific references; indeed, as discussed further below, the sequence of Claim 1’s description would make such references nonsensical. As the Federal Circuit has observed, in situations involving two arguably divergent independent claims, the principle of “claim differentiation takes on relevance in the context of a claim construction that would render additional, or different, language in another independent claim superfluous.” *Curtiss-Wright Flow Control Corp. v. Velan, Inc.*, 438 F.3d 1374, 1381 (Fed. Cir. 2006). Adopting Defendant’s proposed definition would have precisely that effect with respect to independent Claims 8 and 11.

Third, it appears that Defendant’s proposed definition would not even be factually accurate with respect to certain uses of this term in the claims. For example, looking at the language and structure of independent Claim 1 of the ‘807 patent, the “obtaining” happens *prior to* the repeater selector mechanism’s identification of a repeater server to handle the client request. *See* ‘807 patent col. 26:37–42. Accordingly, it cannot be the case that “obtain[ing]” in Claim 1 means “receiving a client request directed by the repeater selector mechanism to a

repeater server,” because the “obtained” client request has not yet been directed, or even analyzed, by the repeater selector mechanism, and it is not even necessarily a repeater server that is doing the “obtaining.”⁴ Using Defendant’s proposed definition of this term in Claim 1 would render the next step after the term is used (*id. col.* 26:40–42) redundant.

c. **Definition (“Obtain[ing]”)**

The Court adopts no definition because none is needed; both the term “obtain[ing]” and the related phrase at issue have plain and ordinary meanings.

13. *Determining . . . Based . . . On a Name by Which the Repeater Server (Is) Addressed*

a. **Proposed Definitions**

Plaintiff: No proposed construction; plain and ordinary meaning.

Defendant: “Us(es)/(ing) a table to determine the origin server name associated with the name by which the repeater server was addressed, and then determin(es)/(ing) if the origin server belongs to a subscriber.”

b. **Discussion**

Plaintiff maintains that no construction is necessary because the term, as used in the patents at issue, has a plain and ordinary meaning. Defendant, on the other hand, offers a

⁴ One might also take issue more generally with the use of the verb “obtain[ing]” in this way, at all. “Obtaining,” by definition, connotes action on the part of the obtainer to get that which is obtained. One need look no further for an example of this than the ‘807 patent’s use of “obtaining” elsewhere in independent Claim 11 and in dependent claims 3, 4, 14, 15, 23, and 24. In the context of the claims at issue with respect to this term, though, it is clearly the client that is generating and sending the “obtained” request; the repeater selector mechanism or repeater server, as the case may be, merely receives or intercepts that client request—it does not actively “go out and get” that request in any sense. Thus, for that purpose, a more passive verb such as “receiv[ing]” might have been more appropriate.

definition effectively limiting the basis for this determination to a table cross-referencing origin server and repeater server names.

Although Defendant is correct to point out that the ‘807 patent’s specification states that the “repeater uses an internal table to verify that the origin server belongs to a known subscriber,” the Court notes that Defendant’s other textual reference (‘807 patent col. 10:1–8) is to a description of an alternate approach, and thus not necessarily controlling with respect to the primary approach contemplated by the specification text. Moreover, the Court notes that this phrase always appears in tandem in the claim language with the phrase “using at least the subscriber verifying mechanism” or a variant thereof. When placed in context, then, it becomes clear that this determination is made *by* “at least” that mechanism *using* “at least” the name by which the repeater server is addressed. This fact is particularly salient in the formulation used in independent Claim 11:

[D]etermining, using at least the subscriber verifying mechanism, whether the requested information is from a subscriber of the plurality of subscribers, *said determination being based, at least in part, on* at least one of (a) a name by which the repeater server was addressed, and (b) an origin server name in a Uniform Resource Locator (URL) used to make the client request

‘807 patent col. 28:51–57. This phrase is merely an aspect of the operation of the subscriber verifying mechanism, and is properly dealt with in connection with the Court’s analysis of that term. Accordingly, it is not necessary to provide a separate construction of this term. Moreover, the Court believes that this phrase, when read in context, is essentially self-explanatory. Accordingly, the Court agrees with Plaintiff that no construction is necessary for this term.

c. Definition

The Court adopts no definition because none is needed; both the term “determine” and the related phrase at issue have plain and ordinary meanings.

C. Claim Constructions Still in Dispute in the ‘405 Patent

The Court will address each of the terms in the ‘405 patent still in dispute in the order in which those terms are listed in Exhibit A to the parties’ Joint Statement. The claims of the ‘405 patent are organized as follows: Claims 1 and 4 are independent in nature, and respectively claim a “method” and a “computer readable medium comprising computer code” to gather and use “real-time traffic information” to evaluate various possible network paths along which a particular data packet could travel and then select, on the basis of that information, the optimal network path. Claims 2 and 5, which are respectively dependent upon Claims 1 and 4, respectively claim a “method” and a “computer readable medium . . . further comprising computer code” to transmit that data packet along the selected alternate path from the dynamic router to the data packet’s destination. Claims 3 and 6, which are also respectively dependent upon Claims 1 and 4, respectively claim a “method” and a “computer readable medium” “wherein the alternate path comprises an overlay node.”

Claims 7 and 10 are independent in nature, and respectively claim a “method” and an “apparatus” for dynamically transmitting data packets from a router to a destination along a multi-“hop” path, gathering and using “real-time traffic information” at each router “hop” in the data packet’s path to evaluate various possible next router “hops” in the data packet’s path and select the optimal next router “hop.” Claims 8 and 11, which are respectively dependent upon Claims 7 and 10, respectively claim a “method” and “apparatus” in which a routing table entry

associated with a data packet's destination is ignored in selecting the next router "hop." Claims 9 and 12, which are also respectively dependent upon Claims 7 and 10, respectively claim a "method" and "apparatus" in which a routing table entry associated with a data packet's destination is modified in selecting the next router "hop."

As a preliminary matter, the Court notes that the parties' proposed definitions divergently use "message" and "data packet" to describe what appears to be a similar, if not entirely identical, concept. Since the '405 patent primarily uses the phrase "data packet" in the specification and claims, the Court will also use that phrase in its discussions and definitions below, without prejudice to Plaintiff's use of the term "message."

1. Alternative/Alternate Path

a. Proposed Definitions

Plaintiff: "A path that is not the default path."

Defendant: "Path that includes intermediate nodes that are co-located with normal nodes of the communications network, but have additional processing capabilities to form a virtual topology on the top of the existing network."

b. Discussion

This term figures in all of the independent claims in the '405 patent, as well as in dependent Claims 2, 3, 5, and 6. As suggested by Plaintiff's proposed definition (and the everyday meaning of the constituent words), this term is used in contradistinction to "default path." Plaintiff's proposed definition of this term is terse almost to the point of being tautological; *i.e.*, technically correct but insufficiently descriptive. Defendant's construction, on

the other hand, arguably overreaches, attempting to encapsulate several distinct concepts from the ‘405 patent’s invention within the definition of this single term.

Plaintiff argues that the Court should adopt its construction of this term in part because it was the agreed-upon construction in a prior litigation involving an alleged “parent” of the ‘405 patent at issue here. *See Cable & Wireless Internet Servs., Inc. v. Akamai Techs., Inc.*, 272 F. Supp. 2d 912 (N.D. Cal. 2003). The fact that Plaintiff’s proposed construction here was agreed upon by parties to an unrelated litigation that are non-parties here⁵ is useful, but certainly not dispositive. There, the litigants were able to agree on a definition of this term as used in a different (albeit related) patent and within the context of constructions of various other terms in that patent. Accordingly, this aspect of Plaintiff’s arguments is not dispositive of this issue.

Defendant’s proposed construction of this term, however, cannot prevail for at least two reasons. First, Defendant’s construction would effectively limit the invention to the embodiment described in dependent Claims 3 and 6 of the ‘405 patent, “wherein the alternate path comprises an overlay node.” ‘405 patent col. 13:7–8, 38–39. As noted above, the principle of claim differentiation counsels that “the presence of a dependent claim that adds a particular limitation gives rise to a presumption that the limitation in question is not present in the independent claim.” *Phillips*, 415 F.3d at 1315. To construe this term as Defendant proposes would render dependent Claims 3 and 6 superfluous.

⁵ Plaintiff here is the successor-in-interest, through several successive acquisitions, to the plaintiff in the cited litigation with respect to the patents at issue in this case. *See Complaint ¶¶ 27–32*. Akamai Technologies, Inc. is a competitor of both Plaintiff and Defendant and a non-party to the instant case.

Defendant's proposed definition is also inappropriate because it is overbroad. It attempts to capture several aspects of the invention in the definition of only one component of that invention. This term is not used in a vacuum in the patent claims and specification, but instead in conjunction with other terms and concepts, such as "default path" (the construction of which the parties have agreed upon), "a special group of intermediate nodes" ('405 patent col. 2:47), and "intermediate overlay nodes" (*id.* col. 2:55–56). Indeed, the abstract describes "an alternate data forwarding path through one or more overlay nodes." '405 patent at 1. Similarly, before discussing the specific "overlay network" preferred embodiment, the specification's general summary of the invention describes "alternative paths" as "passing through a special group of intermediate nodes." *Id.* col. 2:43–49. The summary further describes messages being transmitted along alternate paths by passing through "intermediate overlay nodes." *Id.* col. 2:52–58. Defendant's proposed definition, which itself appears to contain not only the definition of the words "alternate/alternative" and "path," but also a comprehensive definition of what an intermediate or overlay node is, might conceivably be appropriate if "overlay node" were not separately defined. But that is not the case here: "overlay node" is, in fact, another term in dispute in this litigation, which the Court construes separately below. In light of that, there is no need to incorporate those details about the nature of an "overlay node" into the definition of this term.

Accordingly, Defendant's proposed construction is needlessly overbroad, and will not be adopted. In light of the foregoing, Plaintiff's definition, though terse, is adequate, and will be adopted by the Court. This definition, of course, will not preclude Defendant from arguing that the alternate paths contemplated by the '405 patent necessarily involve special intermediate

nodes. This fact even appears to be acknowledged by Plaintiff. *See* Opening Mem. at 23.

Defendant will simply have to make that argument on the basis of the context in which this term is used, not of the definition of the term itself.

c. Definition

“A path that is not the default path.”

2. *Destination*

a. Proposed Definitions

Plaintiff: “A network connection point at which a message is received.”

Defendant: “The endpoint node identified by the destination field of a data packet.”

b. Discussion

This term appears, either by itself or as part of the phrase “destination field,” in all of the independent claims, as well as some of the dependent claims, of the ‘405 patent. In the latter form, it appears in tandem with the phrase “source field.” This juxtaposition is useful in determining the term’s intended meaning. The source field contains information identifying the origination point of a given data packet; by contradistinction, the destination field contains information identifying the terminus of that data packet. Accordingly, contrary to Plaintiff’s proposed construction, and as Defendant argues, “destination” cannot simply mean any network connection point at which a message is received, but rather must connote the intended *final* network connection point. Much in the same vein, Plaintiff’s proposed definition is problematic because it might encompass not only the intended final destination of a data packet, but also the individual “hops” in a multi-hop path contemplated by Claims 7–12.

The fact that “destination” is used without any article or with the article “a” at the outset of certain claims but then appears in subsequent portions of those same claims preceded by the article “the” further supports this conclusion. Although this distinction might seem minor or overly pedantic at first glance, the Court considers it to be meaningful. The former use communicates the obvious fact that “destination” as contemplated by this patent does not have a fixed meaning associated with one particular point on the Internet, but can instead be any intended destination of a data packet (*e.g.*, any URL, IP address, or other network or Internet location); hence, *a* destination. The latter use within the same claims, on the other hand, makes clear that a *given* data packet will only have one destination; hence, *the* destination for that *particular* data packet.

Use of this term in the specification is not to the contrary. The summary and the detailed description of the invention discuss “source and destination points” (*see, e.g.*, ‘405 patent cols. 2:28, 3:36), “source and destination endpoints” (*id.* col. 9:64) and transmitting data packets “from the source . . . to the destination” (*id.* col. 2:45–47). As in the claims, “destination” is clearly used throughout the specification to refer to the intended terminus for a given data packet, not merely to any “network connection point at which a message is received.” Therefore, Plaintiff’s proposed definition is insufficient.

As for Defendant’s proposed definition, although it accurately captures the terminal character of the term, it also is flawed, in that it arguably creates a problem of recursion: namely, the definition itself contains the very word it is trying to define, as part of the phrase “destination field.” Thus, the definition, even if correct, is insufficiently descriptive.

In light of the foregoing, the Court has decided not to adopt either party's proposed definition, but instead to adopt what it considers to be a compromise definition, which addresses the legitimate arguments made by both parties in connection with this term. The Court's definition emphasizes the terminal character of the claim term—*i.e.*, that it denotes the intended *final* destination of a data packet, not just any connection point through which the data packet may pass on its way to that final destination—but also avoids the arguably recursive nature of Defendant's proposed definition.

c. **Definition**

“The endpoint to which a data packet is to be ultimately transmitted.”

3. *Overlay Node*

a. **Proposed Definitions**

Plaintiff: “A network connection point that has additional functionality for exploiting overlay routing, and that cooperates to provide forwarding to paths overlaid over the underlying network.”

Defendant: “A node of an overlay network that has additional functionality for measuring the cost of communication to all other overlay nodes and that cooperates to provide paths utilizing the underlying network.”

b. **Discussion**

The parties' competing definitions share several concepts: a point in a “network . . . that has additional functionality . . . that cooperates to provide . . . paths” that are related to “the underlying network.” Accordingly, the Court has endeavored to incorporate these agreed aspects into its definition. The parties' definitions diverge, however, on the fundamental question of

whether the “additional functionality” of the overlay nodes necessarily involves the nodes measuring the cost to all, or only some, of the other overlay nodes in the network.

Plaintiff’s definition is taken word-for-word from the definition of “intermediate node” proposed by the plaintiff (as noted above, this Plaintiff’s predecessor-in-interest), and adopted by the court, in *Cable & Wireless*, 272 F. Supp. 2d 912. Although this definition is certainly a relevant and useful resource for the instant question, this Court agrees with Defendant that Plaintiff’s proposed definition is arguably too broad, because the specification of the patent at issue in this case appears to identify overlay nodes not as entirely synonymous with, but rather as a type or subset of, intermediate nodes. *See* ‘405 patent col. 2:43–49. Defendant’s point that “overlay node” is only used in the specification in the context of a particular preferred embodiment involving an “overlay network” is also well taken. However, that does not necessarily mean that this limitation has to be embodied within the definition of “overlay node” itself. The construction of this term that the Court has chosen to adopt will not preclude Defendant from arguing this point, but it also will not do Defendant’s job for it: again, Defendant will have to make its argument on the basis of this term’s context in the specification and claims, and not merely rely on an expansive definition of the term itself to win Defendant’s battle for it.

c. Definition

“Node that is part of a network that has additional functionality to participate in the discovery and cost evaluation of alternate paths; further, the node cooperates to provide forwarding of a received data packet towards a destination using the destination field in the received data packet through paths overlaid over the underlying network.”

4. ***Dynamic Router***

a. **Proposed Definitions**

Plaintiff: “A router capable of transmitting data along a default or alternative path to an intended destination based on real-time traffic information.”

Defendant: “A router that computes on-demand alternative paths for a specific packet to the intended destination, based on real-time traffic information.”

b. **Discussion**

The phrase “dynamic router” appears in independent Claims 1 and 4 and dependent claims 2 and 5. However, although the dynamic nature of the invention is discussed at length in the specification, this term itself notably does not appear anywhere in the specification.

The parties’ disagreement here focuses on whether the dynamic router itself must be the device that performs the computations about alternate paths contemplated in Defendant’s proposed definition. Plaintiff correctly points out that the language of the specification allows for the possibility that devices other than the dynamic router itself can perform those computations and broadcast the results to the dynamic router. *See* ‘405 patent col. 5:37–6:8. However, Defendant argues correctly that Plaintiff’s proposed definition is insufficient, because it fails to describe what is dynamic about the router. Its dynamism stems not from the fact that such a router performs these calculations, but rather from the fact that it *uses* the results of those calculations to select the optimal path for a particular data packet going to a particular destination at a particular moment and then transmits the data packet along that selected path towards its destination. Accordingly, the Court has modified Plaintiff’s proposed definition to reflect this aspect.

c. **Definition**

“A router capable of selecting a default path or an alternate path to a data packet’s intended destination based on real-time traffic information and transmitting that data packet along the selected path.”

5. *Real-time Traffic Information*

a. **Proposed Definitions**

Plaintiff: “Current traffic information at the time of measurement.”

Defendant: “Data concerning the load carried by links or channels of the communications network within a few seconds of the time such data is requested.”

b. **Discussion**

This term appears in all four of the independent claims of the ‘405 patent, but appears nowhere in the specification. The parties’ primary dispute with respect to this term is whether the term “real-time” necessarily implies a concrete time period. Defendant seeks to limit the concept of “real-time” to “within a few seconds” of the request. Plaintiff argues that the concept, as contemplated by the patent, need not have any such concrete limitation. Defendant also argues that Plaintiff’s proposed definition does not define “traffic” in this context.

At the outset, the Court notes that all of the constituent words in this term are used in everyday parlance. The Court does not believe that any of these words means anything other than its respective ordinary meaning, and nothing in the claims or the specification suggests otherwise. Moreover, the Court believes that an average juror would be able to apply his or her everyday understanding of automobile traffic, for example, and intuitive concepts of basic physics (e.g., water flowing faster through an unclogged drain than through a partially clogged

one) to come to a sufficient understanding of “traffic” in the context of computer-based network telecommunications.

As to the parties’ primary dispute with respect to this term, here, again, Defendant has identified a valid aspect of the term’s meaning that is arguably omitted from Plaintiff’s definition (the element of time and its significance here), but has itself failed to articulate that aspect adequately in its own proposed definition. A fundamental goal of this invention is to increase speed and efficiency in network communications between computers by creating a mechanism to decide, on a case-by-case basis and using current network traffic information, which of various paths will be optimal for each particular transmission at the particular time of that transmission. To accomplish this, of course, the network traffic information that a dynamic router uses to make this decision must still be current enough at the actual moment of path selection to portray accurately which path is, in fact, the optimal one. Moreover, in order for this invention to be useful in actually increasing speed or network performance, the delay between the moment of path selection and the moment of actual transmission along the selected path cannot be so great that the selected path is no longer the optimal one when the transmission actually occurs. Although the Court agrees with Plaintiff that it is not necessary to link “real-time” to a concrete time period based on the language of the patent, as a practical matter, in this context, real-time does probably connote something on the order of milliseconds or seconds. The speed of modern telecommunications is such that the load on a particular network path might be in a state of constant flux: a path might be virtually free at one moment and, in a matter of seconds, heavily laden with transmissions. Defendant’s proposed definition is therefore likely *de facto* correct,

but nevertheless legally unnecessary. Accordingly, the Court has decided to adopt, without unnecessary elaboration, a slight variation on Plaintiff's definition.

c. Definition

"Current information about network traffic."

IV. CONCLUSION

For the reasons set forth above, the Court issues this Opinion and Order as the construction of the disputed claim terms in the '935, '807, and '405 patents.

The Clerk is **REQUESTED** to send a copy of this Opinion and Order to counsel of record for the parties.

It is so **ORDERED**.

/s/ MJD

Mark S. Davis
UNITED STATES DISTRICT JUDGE

Norfolk, Virginia
December 10, 2008